

CLAIMS

Having thus described our invention, what we claim as new, and desire to secure Letters Patent is:

1. A method for classifying a data packet in a data processing device to determine which of a plurality of predefined processing rules applies to the data packet, each said rule being associated with a range of possible data values in each of a plurality of dimensions corresponding to respective data items in the data packet, the method comprising:
 - (a) for each dimension, identifying the corresponding data value in the packet and determining which of a set of predefined basic ranges contains that data value, wherein the basic ranges correspond to respective non-overlapping value ranges between successive rule range boundaries in that dimension;
 - (b) for the basic range so determined for each dimension, selecting a corresponding basic range identifier from a set of predefined basic range identifiers corresponding to respective basic ranges in that dimension, wherein, for each of at least two dimensions, the basic range identifiers comprise respective p_D -bit strings generated independently for that dimension by
 - 15 - (b1) defining a hierarchy of primitive ranges, which correspond to respective portions of the rule ranges between rule range boundaries in that dimension such that each rule range in the dimension is represented by one or more primitive ranges, wherein primitive ranges in the same level of the hierarchy are non-overlapping and each primitive range in a level above the lowest hierarchy level is a subset of a primitive range in the level below,
 - 20 - (b2) assigning a primitive range identifier to each primitive range in the hierarchy such that primitive ranges in the lowest level have different identifiers and, in each higher level, any ranges which are subsets of the same range in the level below have different identifiers, and
 - (b3) producing a unique p_D -bit basic range identifier for each basic range such that the p_D -bit identifier comprises a concatenation of the primitive range identifiers for the primitive ranges intersected by that basic range, the primitive range identifiers being concatenated in order of increasing hierarchy level;
 - (c) combining the basic range identifiers for said plurality of dimensions to produce a search key;
 - (d) supplying the search key to a ternary content-addressable memory; and

(e) in the memory, comparing the search key with a set of prestored ternary rule vectors, each associated with a said rule and derived for that rule in dependence on the hierarchy for each of said at least two dimensions, to identify at least one rule which applies to the data packet.

5 2. A method according to claim 1 wherein, for each of said plurality of dimensions, the basic range identifiers comprise respective p_D -bit strings generated independently for that dimension by steps (b1) to (b3).

3. A method according to claim 1 wherein, in step (b3), each p_D -bit basic range identifier begins 10 with said concatenation of primitive range identifiers, and, in step (c), the search key is produced by concatenating the basic range identifiers for said plurality of dimensions.

4. A method according to claim 3 wherein each ternary rule vector is generated from a concatenation of subvectors, corresponding to respective said dimensions, for the associated rule, 15 where, for each of said at least two dimensions, for each rule a said subvector corresponding to each of the primitive ranges representing the rule range in that dimension is defined such that the subvector corresponding to a primitive range begins with a concatenation, in order of increasing hierarchy level, of the primitive range identifier of that range with the primitive range identifier of any lower-level primitive range of which that range is a subset, the resulting concatenation being made up to p_D -bits 20 by bits of value X.

5. A method according to claim 4 wherein the set of rule vectors comprises, for each rule, a rule vector, comprising said concatenation of subvectors, for each possible combination of subvectors for that rule.

25

6. A method according to claim 4 wherein the set of rule vectors comprises, for each rule, a reduced set of rule vectors obtained by eliminating redundancy in the set of all possible combinations of subvectors for that rule in a concatenation of the subvectors corresponding to respective said dimensions.

7. A method according to claim 1 wherein the primitive range identifiers assigned in step (b2) are such that the identifiers assigned to primitive ranges in the lowest hierarchy level are prefix unique, and the identifiers assigned to any ranges which are subsets of the same range in the level 5 below are prefix unique.

8. A method according to claim 1 wherein the primitive range identifiers assigned in step (b2) are such that primitive range identifiers in the same hierarchy level have the same number of bits.

10 9. A method according to claim 1 wherein a priority order is defined for said plurality of rules, and wherein step (e) includes selecting the highest priority rule which applies to the data packet.

10. A method according to claim 1, the method including, prior to performing step (a) for a first data packet:

15 generating said set of basic range identifiers for each of said plurality of dimensions; and deriving said set of ternary rule vectors in dependence on the hierarchy for each of said at least two dimensions, and storing the rule vectors in the ternary content-addressable memory.

11. Apparatus for classifying a data packet to determine which of a plurality of predefined 20 processing rules applies to the data packet, each said rule being associated with a range of possible data values in each of a plurality of dimensions corresponding to respective data items in the data packet, the apparatus comprising:

a first memory storing, for each said dimension, a set of basic range identifiers corresponding to 25 respective basic ranges in that dimension, wherein the basic ranges correspond to respective non-overlapping value ranges between successive rule range boundaries in that dimension, and for each of at least two dimensions, the basic range identifiers comprise respective p_D -bit strings generated independently for that dimension by

- defining a hierarchy of primitive ranges, which correspond to respective portions of the rule ranges between rule range boundaries in that dimension such that each rule range in the

dimension is represented by one or more primitive ranges, wherein primitive ranges in the same level of the hierarchy are non-overlapping and each primitive range in a level above the lowest hierarchy level is a subset of a primitive range in the level below,

- assigning a primitive range identifier to each primitive range in the hierarchy such that primitive ranges in the lowest level have different identifiers and, in each higher level, any ranges which are subsets of the same range in the level below have different identifiers, and
- producing a unique p_D -bit basic range identifier for each basic range such that the p_D -bit identifier comprises a concatenation of the primitive range identifiers for the primitive ranges intersected by that basic range, the primitive range identifiers being concatenated in order of increasing hierarchy level;

10 a ternary content-addressable memory storing a set of ternary rule vectors, each associated with a said rule and derived for that rule in dependence on the hierarchy for each of said at least two dimensions; and

15 control logic configured to identify the data value in the packet corresponding to each said dimension, to access the first memory to retrieve the basic range identifier corresponding to the basic range containing the data value so identified for each dimension, to combine the basic range identifiers for said plurality of dimensions to produce a search key, and to supply the search key to the ternary content-addressable memory for comparison with said ternary rule vectors to identify at least one rule which applies to the data packet.

20 12. Apparatus according to claim 11 wherein the control logic is further configured such that, prior to classifying a first data packet, the control logic generates said set of basic range identifiers for each of said plurality of dimensions and stores the basic range identifiers in the first memory, and derives said set of ternary rule vectors in dependence on the hierarchy for each of said at least two dimensions and stores the rule vectors in the ternary content-addressable memory.

25 13. A data processing device for connection as a node of a network system, the device comprising:

communications circuitry for communication of data packets between the device and other nodes of the network system; and

packet classification apparatus according to claim 11, wherein the control logic is connected to said communications circuitry for receiving data packets to be classified therefrom.

5

14. A computer program product comprising a computer readable medium having embodied therein computer readable program code means for causing a processor of a data processing device to perform a method for classifying data packets to determine which of a plurality of predefined processing rules applies to each data packet, each said rule being associated with a range of possible 10 data values in each of a plurality of dimensions corresponding to respective data items in the data packet, the method comprising:

(a) for each dimension, identifying the corresponding data value in the packet and determining which of a set of predefined basic ranges contains that data value, wherein the basic ranges correspond to respective non-overlapping value ranges between successive rule range boundaries in that dimension;
15 (b) for the basic range so determined for each dimension, selecting a corresponding basic range identifier from a set of predefined basic range identifiers corresponding to respective basic ranges in that dimension, wherein, for each of at least two dimensions, the basic range identifiers comprise respective p_D -bit strings generated independently for that dimension by

20 - (b1) defining a hierarchy of primitive ranges, which correspond to respective portions of the rule ranges between rule range boundaries in that dimension such that each rule range in the dimension is represented by one or more primitive ranges, wherein primitive ranges in the same level of the hierarchy are non-overlapping and each primitive range in a level above the lowest hierarchy level is a subset of a primitive range in the level below,
- (b2) assigning a primitive range identifier to each primitive range in the hierarchy such that primitive ranges in the lowest level have different identifiers and, in each higher level, any ranges which are subsets of the same range in the level below have different identifiers, and
25 - (b3) producing a unique p_D -bit basic range identifier for each basic range such that the p_D -bit identifier comprises a concatenation of the primitive range identifiers for the primitive ranges

intersected by that basic range, the primitive range identifiers being concatenated in order of increasing hierarchy level;

- (c) combining the basic range identifiers for said plurality of dimensions to produce a search key;
 - (d) supplying the search key to a ternary content-addressable memory;
- 5 (e) in the memory, comparing the search key with a set of prestored ternary rule vectors, each associated with a said rule and derived for that rule in dependence on the hierarchy for each of said at least two dimensions, to identify at least one rule which applies to the data packet; and
- (f) prior to performing step (a) for a first data packet, generating said set of basic range identifiers for each of said plurality of dimensions, deriving said set of ternary rule vectors in dependence on the
10 hierarchy for each of said at least two dimensions, and storing the rule vectors in the ternary content-addressable memory.

15. A method for classifying data packets in a data processing device according to the values in respective data packets of a predetermined data item in the data packet format, a plurality of classification results being defined in the device for respective ranges of values of the data item,
15 wherein the method comprises for each data packet:

- (a) segmenting said predetermined data item into a plurality of segments of predetermined length;
- (b) for each segment, selecting, in dependence on the value of that segment, a corresponding segment identifier from a predefined set of segment identifiers associated with respective ranges of segment values, the corresponding identifier being that associated with a range containing the value of said segment;
- (c) combining the segment identifiers selected for said plurality of segments to produce a search key;
- 25 (d) supplying the search key to a ternary content-addressable memory; and
- (e) in the memory, comparing the search key with a set of prestored ternary classification vectors, each associated with a said classification result and derived for that classification result in dependence on the segment identifiers corresponding to segment values contained in the range of data item values for that classification result, to identify a classification result for the data packet.

16. A method according to claim 15 wherein said predetermined data item is an address field of the data packet format such that the value of the address field indicates a network address in a data communications network system in which said processing device is connected in use.

5

17. A method according to claim 16 wherein the address field is a destination address field and wherein each said classification result indicates a next hop in the network system for forwarding of data packets by the data processing device.

10 18. A method according to claim 15 wherein:

for the range of data item values corresponding to each classification result, a classification range is defined in each of a plurality of dimensions corresponding to respective said segments of the data item, the classification range for each dimension corresponding to the range of values of the corresponding segment which are contained in the range of data item values for that classification

15 result;

the set of segment identifiers utilized in step (b) for each segment is predefined for the dimension corresponding to that segment, the segment identifiers being associated with respective basic ranges in that dimension where the basic ranges correspond to respective non-overlapping segment-value ranges between successive classification range boundaries in that dimension; and

20 for at least one said dimension, the segment identifiers comprise respective p_D -bit strings generated by

- (b1) defining a hierarchy of primitive ranges, which correspond to respective portions of the classification ranges between classification range boundaries in that dimension such that each classification range in the dimension is represented by one or more primitive ranges, wherein primitive ranges in the same level of the hierarchy are non-overlapping and each primitive range in a level above the lowest hierarchy level is a subset of a primitive range in the level below,

25

- (b2) assigning a primitive range identifier to each primitive range in the hierarchy such that primitive ranges in the lowest level have different identifiers and, in each higher level, any ranges which are subsets of the same range in the level below have different identifiers, and
- (b3) producing a unique p_D -bit segment identifier for each basic range such that the p_D -bit identifier comprises a concatenation of the primitive range identifiers for the primitive ranges intersected by that basic range, the primitive range identifiers being concatenated in order of increasing hierarchy level;

and wherein each said ternary classification vector is derived for the associated classification result in dependence on the hierarchy for said at least one dimension.

10

19. A method according to claim 18 wherein, for each of at least two said dimensions, the segment identifiers comprise respective p_D -bit strings generated independently for that dimension by steps (b1) to (b3).

15 20. A method according to claim 19 wherein, for each of said plurality of dimensions, the segment identifiers comprise respective p_D -bit strings generated independently for that dimension by steps (b1) to (b3).

21. A method according to claim 18 wherein, in step (b3), each p_D -bit segment identifier begins 20 with said concatenation of primitive range identifiers, and, in step (c), the search key is produced by concatenating the segment identifiers for said plurality of segments.

22. A method according to claim 21 wherein each ternary classification vector is generated from a concatenation of subvectors, corresponding to respective said dimensions, for the associated 25 classification result, where, for said at least one dimension, for each classification result a said subvector corresponding to each of the primitive ranges representing the classification range in that dimension is defined such that the subvector corresponding to a primitive range begins with a concatenation, in order of increasing hierarchy level, of the primitive range identifier of that range

with the primitive range identifier of any lower-level primitive range of which that range is a subset, the resulting concatenation being made up to p_D -bits by bits of value X.

23. A method according to claim 22 wherein the set of classification vectors comprises, for each
5 classification result, a classification vector, comprising said concatenation of subvectors, for each possible combination of subvectors for that classification result.

24. A method according to claim 22 wherein the set of classification vectors comprises, for each classification result, a reduced set of classification vectors obtained by eliminating redundancy in the
10 set of all possible combinations of subvectors for that classification result in a concatenation of the subvectors corresponding to respective said dimensions.

25. A method according to claim 18 wherein the primitive range identifiers assigned in step (b2) are such that the identifiers assigned to primitive ranges in the lowest hierarchy level are prefix
15 unique, and the identifiers assigned to any ranges which are subsets of the same range in the level below are prefix unique.

26. A method according to claim 18 wherein the primitive range identifiers assigned in step (b2) are such that primitive range identifiers in the same hierarchy level have the same number of bits.
20

27. A method according to claim 15 wherein a priority order is defined for said plurality of classification results, and wherein step (e) includes selecting the highest priority classification result which applies to the data packet.

25 28. A method according to claim 18, the method including, prior to performing step (a) for a first data packet:

generating said set of segment identifiers; and

deriving said set of ternary classification vectors in dependence on the segment identifiers and storing the classification vectors in the ternary content-addressable memory.

29. Apparatus for classifying data packets according to the values in respective data packets of a predetermined data item in the data packet format, a plurality of classification results being predefined for respective ranges of values of the data item, the apparatus comprising:

5 control logic configured to segment said predetermined data item in each data packet into a plurality of segments of predetermined length;

a first memory for storing a set of segment identifiers corresponding to respective ranges of segment values for each of said plurality of segments; and

10 a ternary content-addressable memory for storing a set of ternary classification vectors, each associated with a said classification result and derived for that classification result in dependence on the segment identifiers corresponding to segment values contained in the range of data item values for that classification result;

15 wherein the control logic is further configured to access the first memory to retrieve, for each said segment, the segment identifier corresponding to the range of segment values containing the value of that segment, to combine the segment identifiers retrieved for said plurality of segments to produce a search key, and to supply the search key to the ternary content-addressable memory for comparison with said ternary classification vectors to identify a classification result for the data packet.

20 30. A data processing device for connection as a node of a network system, the device comprising:

communications circuitry for communication of data packets between the device and other nodes of the network system; and

25 packet classification apparatus according to claim 29, wherein the control logic is connected to said communications circuitry for receiving data packets to be classified therefrom.

31. A computer program product comprising a computer readable medium having embodied therein computer readable program code means for causing a processor of a data processing device to perform a method for classifying data packets according to the values in respective data packets of

a predetermined data item in the data packet format, a plurality of classification results being predefined for respective ranges of values of the data item, the method comprising for each data packet:

(a) segmenting said predetermined data item into a plurality of segments of predetermined

5 length;

(b) for each segment, selecting, in dependence on the value of that segment, a corresponding segment identifier from a predefined set of segment identifiers associated with respective ranges of segment values, the corresponding identifier being that associated with a range containing the value of said segment;

10 (c) combining the segment identifiers selected for said plurality of segments to produce a search key;

(d) supplying the search key to a ternary content-addressable memory; and

15 (e) in the memory, comparing the search key with a set of prestored ternary classification vectors, each associated with a said classification result and derived for that classification result in dependence on the segment identifiers corresponding to segment values contained in the range of data item values for that classification result, to identify a classification result for the data packet.